

Balanced permutations Even-Mansour ciphers

Shoni Gilboa¹, Shay Gueron^{2,3}, and Mridul Nandi⁴

¹ The Open University of Israel, Raanana 43107, Israel

² University of Haifa, Israel

³ Intel Corporation, Israel Development Center, Israel

⁴ Indian Statistical Institute, Kolkata

August 21, 2015

Abstract. The r -rounds Even-Mansour block cipher is a generalization of the well known Even-Mansour block cipher to r iterations. Attacks on this construction were described by Nikolić et al. and Dinur et al., for $r = 2, 3$. These attacks are only marginally better than brute force, but are based on an interesting observation (due to Nikolić et al.): for a “typical” permutation P , the distribution of $P(x) \oplus x$ is not uniform. This naturally raises the following question. Call permutations for which the distribution of $P(x) \oplus x$ is uniform “balanced”. Is there a sufficiently large family of balanced permutations, and what is the security of the resulting Even-Mansour block cipher?

We show how to generate families of balanced permutations from the Luby-Rackoff construction, and use them to define a $2n$ -bit block cipher from the 2-rounds Even-Mansour scheme. We prove that this cipher is indistinguishable from a random permutation of $\{0, 1\}^{2n}$, for any adversary who has oracle access to the public permutations and to an encryption/decryption oracle, as long as the number of queries is $o(2^{n/2})$. As a practical example, we discuss the properties and the performance of a 256-bit block cipher that is based on our construction, and uses AES as the public permutation.

Keywords: Even-Mansour, block-cipher, Luby-Rackoff

Mathematics Subject Classification: 94A60

1 Introduction

The r -rounds Even-Mansour (EM) block cipher, suggested by Bogdanov et al. [2], encrypts an n -bit plaintext m by

$$\text{EM}_{K_0, K_1, \dots, K_r}^{P_1, P_2, \dots, P_r}(m) = P_r(\dots P_2(P_1(m \oplus K_0) \oplus K_1) \dots \oplus K_{r-1}) \oplus K_r, \quad (1)$$

where $K_0, K_1, \dots, K_r \in \{0, 1\}^n$ are secret keys and P_1, P_2, \dots, P_r are publicly known permutations, which are selected uniformly and independently at random, from the set of permutations of $\{0, 1\}^n$. The confidentiality of the EM cipher is achieved even though the permutations P_1, \dots, P_r are made public. For $r = 1$, (1) reduces to the classical Even-Mansour construction [8].

As a practical example, Bogdanov et al. defined the 128-bit block cipher AES², which is an instantiation of the 2-rounds EM cipher where the two public permutations are AES with two publicly known “arbitrary” keys (they chose the binary digits of the constant π). The complexity of the best (meet-in-the-middle) attack they showed uses $2^{129.6}$ cipher revaluations. Consequently, they conjectured that AES² offers 128-bit security.

Understanding the security of the EM cipher has been the topic of extended research. First, Even and Mansour [8] proved, for $r = 1$, that an adversary needs to make $\Omega(2^{n/2})$ oracle queries before he can decrypt a new message with high success probability. Daemen [5] showed that this bound is tight, by demonstrating a chosen-plaintext key-recovery attack after $O(2^{n/2})$ evaluations of P_1 and the encryption oracle. Bogdanov et al. [2], showed, for the r -rounds EM cipher, $r \geq 2$, that an adversary who sees only $O(2^{2n/3})$ chosen plaintext-ciphertext pairs cannot distinguish the encryption oracle from a random permutation of $\{0, 1\}^n$. This result has been recently improved by Chen and Steinberger [3], superseding intermediate progress made by Steinberger [19] and by Lampe, Patarin and Seurin [12]. They showed that for every r , an adversary needs $\Omega(2^{\frac{r}{r+1}n})$ chosen plaintext-ciphertext pairs before he can distinguish the r -rounds EM cipher from a random permutation of $\{0, 1\}^n$. This bound is tight, by Bogdanov et al.’s [2] distinguishing attack after $O(2^{\frac{r}{r+1}n})$ queries.

Nikolić et al. [15] demonstrated a chosen-plaintext key-recovery attack on the single key variant ($K_0 = K_1 = K_2$) of the 2-rounds EM cipher. Subsequently, Dinur et al. [7] produced additional key-recovery attacks on various other EM variants. All the attack in [15] and [7] are only slightly better than a brute force approach. For example, the attack ([7]) on the single key variant of the 2-rounds EM cipher has time complexity $O\left(\frac{\log n}{n} 2^n\right)$, and the attack ([7]) on AES² (with three different keys) has complexity of $2^{126.8}$ (still better than Bogdanov et al. [2], thus enough to invalidate their that AES² has 2^{128} security).

The above attacks are based on the astute observation, made in [15], that for a “typical” permutation P of $\{0, 1\}^n$, the distribution of $P(x) \oplus x$ over uniformly chosen $x \in \{0, 1\}^n$ is not uniform. Currently, this observation yields only weak attacks, but the unveiled asymmetry may have the potential to lead to stronger results.

This motivates the following question. Call a permutation P of $\{0, 1\}^n$ “balanced” if the distribution of $P(x) \oplus x$, over uniformly chosen $x \in \{0, 1\}^n$, is uniform. Can we construct a block cipher based on balanced permutations? We point out that, a priori, it is not even clear that there exists a family of such permutations, that is large enough to support a block cipher construction.

In this work, we show how to generate a large family of balanced permutations of $\{0, 1\}^{2n}$, by observing that a 2-rounds Luby-Rackoff construction with any two identical *permutations* of $\{0, 1\}^n$, always yields a balanced permutation (of $\{0, 1\}^{2n}$). We use these permutations in an EM setup (illustrated in Figure 2, top panel), to construct a block cipher with block size of $2n$ bits. Note that in this EM setup, the permutations P_1, P_2 are not chosen uniformly at random from the set of all permutations of $\{0, 1\}^{2n}$. They are selected from a particular

subset of the permutations of $\{0,1\}^{2n}$, and defined via a random choice of two permutations of $\{0,1\}^n$, as the paper describes.

For the security of the resulting $2n$ bits block cipher, we would ideally like to maintain the security of the EM cipher (on blocks of $2n$ bits). This would be guaranteed if we replaced the random permutation in the EM cipher, with an indifferntiable block cipher (as defined in [13]). However, the balanced permutations we use in the EM construction are 2-rounds Luby-Rackoff permutations, and it was shown in [4] that even the 5-rounds Luby-Rackoff construction does not satisfy indifferntiability. Therefore, it is reasonable to expect weaker security properties in our cipher. Indeed, we demonstrate a distinguishing (not a key recovery) attack that uses $O(2^{n/2})$ queries. On the other hand, we prove that a smaller number of chosen plaintext-ciphertext queries is not enough to distinguish the block cipher from a random permutation of $\{0,1\}^{2n}$.

We comment that the combination of EM and Luby-Rackoff constructions have already been used and analyzed. Gentry and Ramzan [9] showed that the internal permutation of the Even-Mansour construction for $2n$ -bits block size, can be securely replaced by a 4-rounds Luby-Rackoff scheme with public round functions. They proved that the resulting construction is secure up to $O(2^{n/2})$ queries. Lampe and Seurin [11] discuss r -rounds Luby-Rackoff constructions where the round functions are of the form $x \mapsto F_i(K_i \oplus x)$, F_i is a public random function, and K_i is a (secret) round key. For an even number of rounds, this can be seen as a $r/2$ -rounds EM construction, where the permutations are 2-rounds Luby-Rackoff permutations. They show that this construction is secure up to $O(2^{\frac{tn}{t+1}})$ queries, where $t = \lfloor r/3 \rfloor$ for non-adaptive chosen-plaintext adversaries, and $t = \lfloor r/6 \rfloor$ for adaptive chosen-plaintext and ciphertext adversaries. These works bare some similarities to ours, but the new feature in our construction is the emergence of balanced permutations.

The paper is organized as follows. In Section 2 we discuss balanced permutations and balanced permutations EM ciphers. Section 3 provides general background for the security analysis given in Section 4. In Section 5, we demonstrate the distinguishing attack. A practical use of our construction is a 256-bit block cipher is based on AES. Section 6 defines this cipher and discusses its performance characteristics. We conclude with a discussion in Section 7.

2 Balanced permutations and balanced permutations EM ciphers

2.1 Balanced permutations

Definition 1 (Balanced permutation⁵). Let σ be a permutation of $\{0,1\}^n$. Define the function $\tilde{\sigma} : \{0,1\}^n \rightarrow \{0,1\}^n$ by $\tilde{\sigma}(\omega) = \omega \oplus \sigma(\omega)$, for every $\omega \in \{0,1\}^n$. We say that σ is a balanced permutation if $\tilde{\sigma}$ is also a permutation.

⁵ Also known as “orthomorphism” in the mathematical literature

Example 1. Let $A \in M_{n \times n}(\mathbb{Z}_2)$ be a matrix such that both A and $I + A$ are invertible. Define $\pi_A : \mathbb{Z}_2^n \rightarrow \mathbb{Z}_2^n$ by $\pi_A(x) = Ax$. Then π_A is a balanced permutation of $\{0, 1\}^n$. One such matrix is defined by $A_{i,i} = A_{i,i+1} = 1$ for $i = 1, 2, \dots, n-1$, $A_{n,1} = 1$ and $A_{i,j} = 0$ for all other $1 \leq i, j \leq n$.

Example 2. Let a be an element of $GF(2^n)$ such that $a \neq 0, 1$. Identify $GF(2^n)$ with $\{0, 1\}^n$, so that field addition corresponds to bitwise XOR. The field's multiplication is denoted by \times . The function $x \rightarrow a \times x$ is a balanced permutation of $\{0, 1\}^n$. Note that this example is actually a special case of the previous one.

The balanced permutations provided by the above examples are a small family of permutations, and moreover are all linear. We now give a recipe for generating a large family of balanced permutations, by employing the Feistel construction that turns any function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ to a permutation of $\{0, 1\}^{2n}$.

Let us use the following notation. For a string $\omega \in \{0, 1\}^{2n}$, denote the string of its first n bits by $\omega_L \in \{0, 1\}^n$, and the string of its last n bits by $\omega_R \in \{0, 1\}^n$. Denote the concatenation of two strings $\omega_1, \omega_2 \in \{0, 1\}^n$ (in this order) by $\omega_1 * \omega_2 \in \{0, 1\}^{2n}$. We have the following identities:

$$(\omega_1 * \omega_2)_L = \omega_1, \quad (\omega_1 * \omega_2)_R = \omega_2, \quad \omega_L * \omega_R = \omega.$$

Definition 2 (Luby-Rackoff permutations). Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a function. Let $LR[f] : \{0, 1\}^{2n} \rightarrow \{0, 1\}^{2n}$ be the Luby-Rackoff (a.k.a Feistel) permutation

$$LR[f](\omega) := \omega_R * (\omega_L \oplus f(\omega_R)). \quad (2)$$

For every $r \geq 2$ and r functions $f_1, \dots, f_r : \{0, 1\}^n \rightarrow \{0, 1\}^n$, we define the r -rounds Luby-Rackoff permutation to be the composition

$$LR[f_1, \dots, f_r] := LR[f_r] \circ \dots \circ LR[f_1].$$

Since we use here extensively the special case $LR[f, f]$, we denote it by $LR^2[f]$.

The following proposition shows that when f is, itself, a permutation, then $LR^2[f]$ is a balanced permutation.

Proposition 1. Let f be a permutation of $\{0, 1\}^n$. Then, the 2-rounds Luby-Rackoff permutation, $LR^2[f]$, is a balanced permutation of $\{0, 1\}^{2n}$.

Proof. Denote $P := LR^2[f]$. Observe first that

$$\begin{aligned} P(\omega) &= LR^2[f](\omega) = LR[f](LR[f](\omega)) = LR[f](\omega_R * (\omega_L \oplus f(\omega_R))) = \\ &= (\omega_L \oplus f(\omega_R)) * (\omega_R \oplus f(\omega_L \oplus f(\omega_R))). \end{aligned} \quad (3)$$

Therefore,

$$\tilde{P}(\omega) = f(\omega_R) * f(\omega_L \oplus f(\omega_R)).$$

Assume that $x, y \in \{0, 1\}^{2n}$ such that $\tilde{P}(x) = \tilde{P}(y)$, i.e.,

$$f(x_R) * f(x_L \oplus f(x_R)) = f(y_R) * f(y_L \oplus f(y_R))$$

Then, $f(x_R) = f(y_R)$ and $f(x_L \oplus f(x_R)) = f(y_L \oplus f(y_R))$. Since (by assumption) f is one-to-one, $x_R = y_R$ and $x_L \oplus f(x_R) = y_L \oplus f(y_R)$, it follows that $x_L = (x_L \oplus f(x_R)) \oplus f(x_R) = (y_L \oplus f(y_R)) \oplus f(y_R) = y_L$. We established that $\tilde{P}(x) = \tilde{P}(y)$ implies $x = x_L * x_R = y_L * y_R = y$ which concludes the proof.

Figure 1 shows an illustration of 2-rounds Luby-Rackoff (balanced) permutation.

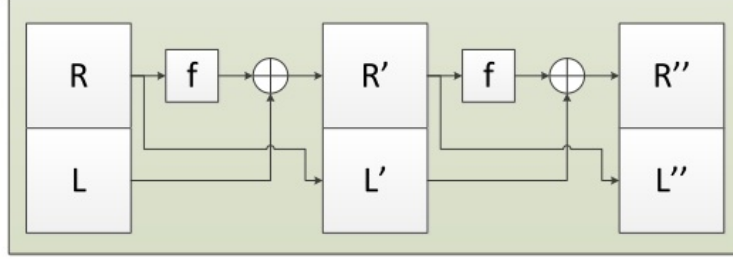


Fig. 1. The figure shows a function from $\{0, 1\}^{2n}$ to $\{0, 1\}^{2n}$, based on two Feistel rounds with a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$. For any function f , this construction is a permutation of $\{0, 1\}^{2n}$, denoted $\text{LR}^2[f]$. We call it a “2-rounds Luby-Rackoff permutation”. Proposition 1 shows that if f itself is a *permutation* of $\{0, 1\}^n$, then $\text{LR}^2[f]$ is a balanced permutation of $\{0, 1\}^{2n}$.

2.2 Balanced permutations EM ciphers

Definition 3 (r -rounds balanced permutations EM ciphers (BPEM)).

Let $n \geq 1$ and $r \geq 1$ be integers. Let K_0, K_1, \dots, K_r be $r + 1$ strings in $\{0, 1\}^{2n}$. Let f_1, f_2, \dots, f_r be r permutations of $\{0, 1\}^n$. Their associated 2-rounds Luby-Rackoff (balanced) permutations (of $\{0, 1\}^{2n}$) are $\text{LR}^2[f_1], \text{LR}^2[f_2], \dots, \text{LR}^2[f_r]$, respectively. The r -rounds balanced permutations EM (BPEM) block cipher is defined as

$$\text{BPEM}[K_0, K_1, \dots, K_r; f_1, \dots, f_r] := \text{EM}_{K_0, K_1, \dots, K_r}^{\text{LR}^2[f_1], \text{LR}^2[f_2], \dots, \text{LR}^2[f_r]}, \quad (4)$$

(where EM is defined by (1)). It encrypts $2n$ -bit blocks with an r -rounds EM cipher with the keys K_0, K_1, \dots, K_r , where the r permutations P_1, P_2, \dots, P_r (of $\{0, 1\}^{2n}$) are set to be $\text{LR}^2[f_1], \text{LR}^2[f_2], \dots, \text{LR}^2[f_r]$, respectively.

The use of the r -rounds BPEM cipher for encryption (and decryption) starts with an initialization step, where the permutations f_1, f_2, \dots, f_r are selected uniformly and independently, at random from the set of permutations of $\{0, 1\}^n$. After they are selected, they can be made public. Subsequently, per session/message, the secret keys K_0, K_1, \dots, K_r are selected uniformly and independently, at random, from $\{0, 1\}^{2n}$. Figure 2 illustrates a 2-rounds BPEM cipher $\text{BPEM}[K_0, K_1, K_2; f_1, f_2]$, which is the focus of this paper.

Remark 1. The r -rounds EM cipher is not necessarily secure with *any* choice of balanced permutations as P_1, P_2, \dots, P_r . For example, it can be easily broken when using the linear balanced permutations shown in Examples 1 and 2.

Remark 2. In our construction, the permutations P_1, P_2, \dots, P_r are not random permutations. Therefore, the security analysis of the “classical” EM does not apply, and the resulting cipher (BPEM) may not be secure. Indeed, it is easy to see that the 1-round BPEM does not provide confidentiality. For any plaintexts $m \in \{0, 1\}^{2n}$, we have, by (3),

$$(\text{LR}^2[f](m \oplus K_0))_L = (m_L \oplus (K_0)_L) \oplus f(m_R \oplus (K_0)_R)$$

Therefore, by (4), (1) and (3),

$$\begin{aligned} (\text{BPEM}[K_0, K_1; f](m))_L &= \left(\text{EM}_{K_0, K_1}^{\text{LR}^2[f]}(m) \right)_L = (\text{LR}^2[f](m \oplus K_0))_L \oplus (K_1)_L = \\ &= m_L \oplus (K_0)_L \oplus (K_1)_L \oplus f(m_R \oplus (K_0)_R). \end{aligned}$$

It follows that if, e.g., $(m_1)_R = (m_2)_R$ then

$$(\text{BPEM}[K_0, K_1; f](m_1) \oplus \text{BPEM}[K_0, K_1; f](m_2))_L = (m_1 \oplus m_2)_L$$

which means that the ciphertexts leak out information on m_1, m_2 . This also implies that the r -rounds BPEM cipher must be used with $r \geq 2$ to have any hope for achieving security.

Remark 3. By construction, $\text{BPEM}[K_0, K_1, \dots, K_r; f_1, \dots, f_r]$ ($r \geq 2$) is immune against any attack that tries to leverage the non-uniformity of $P(x) \oplus x$ (including [15] and [7]). Obviously, this does not guarantee it is secure (as indicated in Remark 1).

In Section 4 we prove that the 2-round BPEM cipher is indistinguishable from a random permutation.

2.3 Equivalent representation of BPEM in terms of LR

In this section we show that 2-rounds BPEM can be viewed as a “keyed”⁶ Luby-Rackoff cipher (in fact, the r -rounds BPEM has a similar representation for every r).

Notation 1 Given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a key $K \in \{0, 1\}^n$ we denote $\text{EM}_{K, K}^f$ by $f^{\oplus K}$, namely

$$\text{EM}_{K, K}^f(x) = f(x \oplus K) \oplus K.$$

⁶ By “keyed” we mean that each function used in the Luby-Rackoff construction is selected from a family of functions indexed by a key.

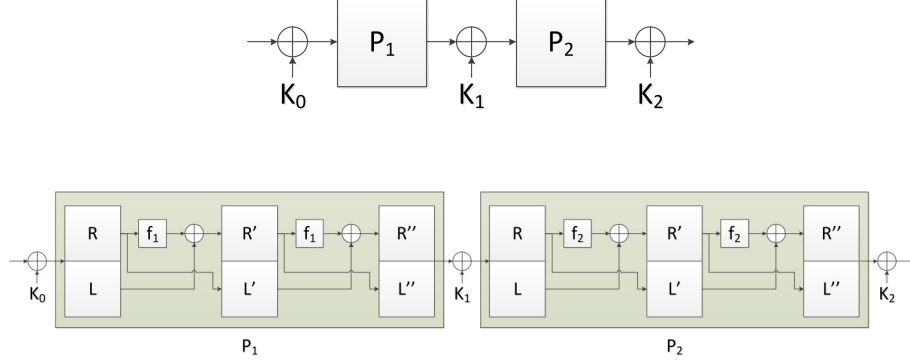


Fig. 2. The 2-rounds balanced permutations EM (BPEM) cipher operates on blocks of size $2n$ bits. The permutations P_1 and P_2 are balanced permutations of $\{0, 1\}^{2n}$, defined as 2-rounds Luby-Rackoff permutations. f_1 and f_2 are two (public) permutations of $\{0, 1\}^n$. Each of K_0, K_1, K_2 is a $2n$ -bit secret key. See explanation in the text.

Lemma 1. Let $K_0, K_1, K_2 \in \{0, 1\}^{2n}$ and let f_1, f_2 be two permutations of $\{0, 1\}^n$. Then,

$$BPEM[K_0, K_1, K_2; f_1, f_2] = LR[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4}] \oplus (K'_6 * K'_5)$$

where

$$\begin{pmatrix} K'_1 \\ K'_2 \\ K'_3 \\ K'_4 \\ K'_5 \\ K'_6 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} (K_0)_R \\ (K_0)_L \\ (K_1)_R \\ (K_1)_L \\ (K_2)_R \\ (K_2)_L \end{pmatrix}. \quad (5)$$

Proof. For every function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, $K \in \{0, 1\}^{2n}$ and $\omega \in \{0, 1\}^{2n}$ we have, by (2),

$$\begin{aligned} LR[f](\omega \oplus K) &= (\omega \oplus K)_R * ((\omega \oplus K)_L \oplus f((\omega \oplus K)_R)) = \\ &= (\omega_R * (\omega_L \oplus f(\omega_R \oplus K_R) \oplus K_R)) \oplus (K_R * (K_L \oplus K_R)) = \\ &= (\omega_R * (\omega_L \oplus f^{\oplus K_R}(\omega_R))) \oplus (K_R * (K_L \oplus K_R)) = \\ &= LR[f^{\oplus K_R}](\omega) \oplus (K_R * (K_L \oplus K_R)) \end{aligned}$$

and hence

$$\begin{aligned} LR^2[f](\omega \oplus K) &= LR[f](LR[f](\omega \oplus K)) = \\ &= LR[f]((LR[f^{\oplus K_R}](\omega)) \oplus (K_R * (K_L \oplus K_R))) = \\ &= LR[f^{\oplus (K_L \oplus K_R)}](LR[f^{\oplus K_R}](\omega)) \oplus ((K_L \oplus K_R) * K_L) = \\ &= LR[f^{\oplus K_R}, f^{\oplus (K_L \oplus K_R)}](\omega) \oplus ((K_L \oplus K_R) * K_L). \end{aligned}$$

In particular

$$\begin{aligned}
\text{LR}^2[f_1](\omega \oplus K_0) &= \\
&= \text{LR} \left[f_1^{\oplus(K_0)_R}, f_1^{\oplus((K_0)_L \oplus (K_0)_R)} \right] (\omega) \oplus (((K_0)_L \oplus (K_0)_R) * (K_0)_L) = \\
&= \text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2} \right] (\omega) \oplus (K'_2 * (K'_1 \oplus K'_2))
\end{aligned}$$

and then

$$\begin{aligned}
\text{LR}^2[f_2] (\text{LR}^2[f_1](\omega \oplus K_0) \oplus K_1) &= \\
&= \text{LR}^2[f_2] \left(\text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2} \right] (\omega) \oplus (K'_2 * (K'_1 \oplus K'_2)) \oplus K_1 \right) = \\
&= \text{LR} \left[f_2^{\oplus(K'_1 \oplus K'_2 \oplus (K_1)_R)}, f_2^{\oplus(K'_1 \oplus (K_1)_L \oplus (K_1)_R)} \right] \left(\text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2} \right] (\omega) \right) \oplus \\
&\quad \oplus ((K'_1 \oplus (K_1)_L \oplus (K_1)_R) * (K'_2 \oplus (K_1)_L)) = \\
&= \text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4} \right] (\omega) \oplus (K'_4 * (K'_3 \oplus K'_4)).
\end{aligned}$$

Therefore, by (4) and (1),

$$\begin{aligned}
\text{BP}^{\text{EM}}[K_0, K_1, K_2; f_1, f_2](\omega) &= \text{EM}_{K_0, K_1, K_2}^{\text{LR}^2[f_1], \text{LR}^2[f_2]}(\omega) = \\
&= \text{LR}^2[f_2] (\text{LR}^2[f_1](\omega \oplus K_0) \oplus K_1) \oplus K_2 = \\
&= \text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4} \right] (\omega) \oplus ((K'_4 \oplus (K_2)_L) * (K'_3 \oplus K'_4 \oplus (K_2)_R)) = \\
&= \text{LR} \left[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4} \right] (\omega) \oplus (K'_6 * K'_5).
\end{aligned}$$

3 Security preliminaries and definitions

Let A be an oracle adversary which interacts with one or more oracles. Suppose that \mathcal{O} and \mathcal{O}' are two oracles (or a tuple of oracles) with same domain and range spaces. We define the distinguishing advantage of A distinguishing \mathcal{O} and \mathcal{O}' as

$$\Delta_A(\mathcal{O}; \mathcal{O}') := |\Pr[A^{\mathcal{O}} = 1] - \Pr[A^{\mathcal{O}'} = 1]|.$$

The maximum advantage $\max_A \Delta_A(\mathcal{O}; \mathcal{O}')$ over all adversaries with complexity θ (which includes query, time complexities etc.) is denoted by $\Delta_\theta(\mathcal{O}; \mathcal{O}')$. When we consider computationally unbounded adversaries (which is done in this paper), the time and memory parameters are not present and so we only consider query complexities. In the case of a single oracle, θ is the number of queries, and in the case of a tuple of oracles, θ would be of the form (q_1, \dots, q_r) where q_i denotes the number of queries to the i^{th} oracle. While we define security advantages of \mathcal{O} , we usually choose \mathcal{O}' to be an ideal candidate, such as the random permutation Π or a random function. The PRP-advantage of A against a keyed construction \mathcal{C}_K is $\Delta_A(\mathcal{C}_K; \Pi)$. The maximum PRP-advantage with query complexity θ is denoted as $\Delta_{\mathcal{C}}^{\text{PRP}}(\theta)$.

In this paper, we always assume that queries to an oracle \mathcal{O} are allowed in both directions, i.e., to \mathcal{O}^{-1} as well. We denote

$$\begin{aligned}\Delta_A^\pm(\mathcal{O}, \mathcal{O}') &:= \Delta_A\left((\mathcal{O}, \mathcal{O}^{-1}); (\mathcal{O}', \mathcal{O}'^{-1})\right), \\ \Delta_\theta^\pm(\mathcal{O}, \mathcal{O}') &:= \Delta_\theta\left((\mathcal{O}, \mathcal{O}^{-1}); (\mathcal{O}', \mathcal{O}'^{-1})\right).\end{aligned}$$

The SPRP-advantage of a keyed construction \mathcal{C}_K (where the adversary has access to both the encryption \mathcal{C}_K and its decryption \mathcal{C}_K^{-1}) is defined by

$$\Delta_{\mathcal{C}}^{\text{sprp}}(\theta) := \Delta_\theta^\pm(\mathcal{C}_K; \Pi).$$

When a construction \mathcal{C} is based on one or more ideal permutations or random permutations f_1, \dots, f_r and a key K , we define SPRP-advantage of a distinguisher A , in presence of ideal candidates, as $\Delta_A^\pm((\mathcal{C}, f_1, \dots, f_r); (\Pi, f_1, \dots, f_r))$ where Π is sampled independently of $\hat{f} := (f_1, \dots, f_r)$. We denote the maximum advantage by $\Delta_{\mathcal{C}}^{\text{im-sprp}}(\theta) := \Delta_\theta^\pm((\mathcal{C}, \hat{f}); (\Pi, \hat{f}))$ which we call SPRP-advantage in the ideal model. The complexity parameters of the above advantages depend on the number of oracles, and will be explicitly declared in specific instances.

We state two simple observations on the distinguishing advantages for oracles (we skip the proofs of these observations, as these are straightforward).

Observation 1 *If $\mathcal{O}_1, \mathcal{O}_2$ and \mathcal{O}' are three independent oracles, then*

$$\Delta_{q,q'}^\pm((\mathcal{O}_1, \mathcal{O}'); (\mathcal{O}_2, \mathcal{O}')) \leq \Delta_q^\pm(\mathcal{O}_1; \mathcal{O}_2).$$

Observation 2 *If \mathcal{C} is an oracle construction, then (by using standard reduction)*

$$\Delta_{q,q'}^\pm((\mathcal{C}^{\mathcal{O}_1}, \mathcal{O}'); (\mathcal{C}^{\mathcal{O}_2}, \mathcal{O}')) \leq \Delta_{r,q,q'}^\pm((\mathcal{O}_1, \mathcal{O}'); (\mathcal{O}_2, \mathcal{O}'))$$

(where r is the number of queries to \mathcal{O} , needed to simulate one query to the construction $\mathcal{C}^{\mathcal{O}}$).

Note that in the Observation 2, we do not need to assume any kind of independence of the oracles. Analogous observations, up to obvious changes, hold for the case where $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}'$ are tuples of oracles.

3.1 Coefficient-H technique

Patarin's coefficient-H technique [16] (see also [17]) is a tool for showing an upper bound for the distinguishing advantage. Here is the basic result of the technique.

Theorem 1 (Patarin [16]). *Let \mathcal{O} and \mathcal{O}' be two oracle algorithms with domain D and range R . Suppose there exist a set $\mathcal{V}_{\text{bad}} \subseteq D^q \times R^q$ and $\varepsilon > 0$ such that the following conditions hold:*

1. *For all $(x_1, \dots, x_q, y_1, \dots, y_q) \notin \mathcal{V}_{\text{bad}}$,*

$$\Pr[\mathcal{O}(x_1) = y_1, \dots, \mathcal{O}(x_q) = y_q] \geq (1 - \varepsilon) \Pr[\mathcal{O}'(x_1) = y_1, \dots, \mathcal{O}'(x_q) = y_q]$$

(the above probabilities are called interpolation probabilities).

2. For all A making at most q queries to \mathcal{O}' , $\Pr[\text{Trans}(A^{\mathcal{O}'}) \in \mathcal{V}_{\text{bad}}] \leq \delta$ where $\text{Trans}(A^{\mathcal{O}'}) = (x_1, \dots, x_q, y_1, \dots, y_q)$, x_i and y_i denote the i^{th} query and response of A to \mathcal{O}' .

Then,

$$\Delta_q(\mathcal{O}; \mathcal{O}') \leq \varepsilon + \delta.$$

The above result can be applied for more than one oracle. In such cases we split the parameter q into (q_1, \dots, q_r) where q_i denotes the maximum number of queries to the i^{th} oracle. Moreover, if we have an oracle \mathcal{O} and its inverse \mathcal{O}^{-1} then the interpolation probability for both \mathcal{O} and \mathcal{O}^{-1} can be simply expressed through the interpolation probability of \mathcal{O} only. For example, if an adversary makes a query y to \mathcal{O}^{-1} and obtains the response x , we can write $\mathcal{O}(x) = y$. Therefore, under the conditions of Theorem 1 we also have $\Delta_q^\pm(\mathcal{O}; \mathcal{O}') \leq \varepsilon + \delta$.

3.2 Known related results

The security of Even-Mansour cipher It is known that the Even-Mansour cipher EM_{K_0, K_1} is SPRP secure in the ideal model, in the following sense: $\Delta_{EM}^{\text{im-sprp}}(q_1, q_2) = O(q_1 q_2 / 2^n)$. The same is true for the single key variant $\text{EM}_{K, K}$. In Section 4, we provide (using Patarin's coefficient-H technique) a simple proof of this result (Lemma 2) and a more general result (Lemma 3).

The security of Luby-Rackoff encryption The 3-rounds Luby-Rackoff construction is PRP secure and the 4-rounds Luby-Rackoff construction is SPRP secure, when the underlying functions f_i are PRP's (or PRF's). We use the following quantified version of the SPRP security of the 4-rounds case.

Theorem 2 (Piret [18]). *Let Π_1, \dots, Π_4 be four independent random permutations of $\{0, 1\}^n$, and let Π be a random permutation of $\{0, 1\}^{2n}$. Then, $\text{LR}[\Pi_1, \dots, \Pi_4]$ is SPRP secure in the following sense:*

$$\Delta_q^\pm(\text{LR}[\Pi_1, \dots, \Pi_4]; \Pi) \leq \frac{5q(q-1)}{2^n}.$$

The above bound $O(q^2/2^n)$ is tight (see [20]). In the proof of Theorem 7 we use the following, more general, result.

Theorem 3 (Nandi [14]). *Let $r \geq 4$, and let $(\alpha_1, \dots, \alpha_r)$ be a sequence of numbers from $\{1, \dots, t\}$ such that $(\alpha_1, \dots, \alpha_r) \neq (\alpha_r, \dots, \alpha_1)$. Let Π_1, \dots, Π_t be t independent random permutations of $\{0, 1\}^n$, and let Π be a random permutation of $\{0, 1\}^{2n}$. Then, $\text{LR}[\Pi_{\alpha_1}, \dots, \Pi_{\alpha_r}]$ is SPRP secure in the following sense:*

$$\Delta_q(\text{LR}[\Pi_{\alpha_1}, \dots, \Pi_{\alpha_r}]; \Pi) \leq \frac{(r^2 + 1)q^2}{2^n - 1} + \frac{q^2}{2^{2n}}.$$

4 Security analysis of our construction

4.1 Security analysis of tuples of single key 1-round EM cipher

Notation 2 Let $x_1, \dots, x_t \in \{0, 1\}^n$. We use $\text{coll}(x_1, \dots, x_t)$ to indicate the existence of a collision, i.e., that $x_i = x_j$ for some $1 \leq i < j \leq t$. Otherwise, we write $\text{dist}_n(x_1, \dots, x_t)$, and say that the tuple (x_1, \dots, x_t) is block-wise distinct. Given a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and a tuple $x_1, \dots, x_t \in \{0, 1\}^n$ we define

$$f^{(t)}(x_1, \dots, x_t) := (f(x_1), \dots, f(x_t)).$$

For positive integers m, r , denote

$$P(m, r) = m(m-1) \cdots (m-r+1).$$

Observation 3 For every $\text{dist}_n(x_1, \dots, x_t)$, $\text{dist}_n(y_1, \dots, y_t)$ and a uniform random permutation Π on $\{0, 1\}^n$,

$$\Pr[\Pi^{(t)}(x_1, \dots, x_t) = (y_1, \dots, y_t)] = \frac{1}{P(2^n, t)}$$

More generally, let Π_1, \dots, Π_r be independent uniform random permutations over $\{0, 1\}^n$ then for every block-wise distinct tuples $X^i, Y^i \in (\{0, 1\}^n)^{t_i}$, $1 \leq i \leq r$ we have

$$\Pr[\Pi_1^{(t_1)}(X^1) = Y^1, \dots, \Pi_r^{(t_r)}(X^r) = Y^r] = \frac{1}{P(2^n, t_1)} \times \cdots \times \frac{1}{P(2^n, t_r)}. \quad (6)$$

Now we show that for a random permutation Π of $\{0, 1\}^n$ and a uniformly chosen K , the permutation $\Pi^{\oplus K}$ (single keyed 1-round EM, see Notation 1) is SPRP secure in the ideal model.

Lemma 2. Let Π and Π_1 be independent random permutations of $\{0, 1\}^n$. Then

$$\Delta_{q_1, q_2}^{\pm}((\Pi^{\oplus K}, \Pi); (\Pi_1, \Pi)) \leq \frac{2q_1 q_2}{2^n}.$$

Proof. We use Patarin's coefficient H-technique. We take the set of bad views \mathcal{V}_{bad} to be the empty set. We need to show that for every tuples $M, C \in (\{0, 1\}^n)^{q_1}$, $x, y \in (\{0, 1\}^n)^{q_2}$,

$$\begin{aligned} & \Pr[\Pi^{\oplus K}(M_i) = C_i, 1 \leq i \leq q_1, \Pi(x_i) = y_i, 1 \leq i \leq q_2] \geq \\ & \geq (1 - \varepsilon) \Pr[\Pi_1(M_i) = C_i, 1 \leq i \leq q_1, \Pi(x_i) = y_i, 1 \leq i \leq q_2], \end{aligned}$$

where $\varepsilon = \frac{2q_1 q_2}{2^n}$. With no loss of generality we may assume that each of the tuples M, C, x, y is block-wise distinct. Then, by (6),

$$\begin{aligned} I_{ideal} &:= \Pr[\Pi_1(M_i) = C_i, 1 \leq i \leq q_1, \Pi(x_i) = y_i, 1 \leq i \leq q_2] = \\ &= \Pr[\Pi_1^{(q_1)}(M) = C, \Pi^{(q_2)}(x) = y] = \frac{1}{P(2^n, q_1)} \times \frac{1}{P(2^n, q_2)}. \end{aligned}$$

We say that a key $K \in \{0, 1\}^n$ is “good” if $K \oplus M_i \neq x_j$ and $K \oplus C_i \neq y_j$ for all $1 \leq i \leq q_1, 1 \leq j \leq q_2$. In other words, for a good key all the inputs (outputs) of Π (in the I_{real} computation) are block-wise distinct. Thus, for any given good key K ,

$$\begin{aligned} \Pr[\Pi(M_i \oplus K) = (K \oplus C_i), 1 \leq i \leq q_1, \Pi(x_j) = y_j, 1 \leq j \leq q_2] &= \\ &= \frac{1}{P(2^n, q_1 + q_2)} \geq I_{ideal}. \end{aligned}$$

By a simple counting argument, the number of good keys is at least $2^n - 2q_1q_2$, i.e., the probability that a randomly chosen key is good, is at least $(1 - \varepsilon)$, where $\varepsilon = \frac{2q_1q_2}{2^n}$. Therefore, we have

$$I_{real} := \Pr[\Pi^{\oplus K}(M_i) = C_i, 1 \leq i \leq q_1, \Pi(x_j) = y_j, 1 \leq j \leq q_2] \geq (1 - \varepsilon)I_{ideal}$$

and the result follows by Theorem 1.

Now, we extend Lemma 2 to a tuple $(\Pi_{\alpha_1}^{\oplus K_{\beta_1}}, \dots, \Pi_{\alpha_t}^{\oplus K_{\beta_t}})$ of single key 1-round EM encryptions, where some keys and permutations can be repeated.

Lemma 3. *Let $\Pi_1, \dots, \Pi_r, \bar{\Pi}_1, \dots, \bar{\Pi}_t$ be independent random permutations of $\{0, 1\}^n$ and K_1, \dots, K_s be chosen uniformly and independently from $\{0, 1\}^n$. We write $\hat{\Pi}$ to denote (Π_1, \dots, Π_r) . Let $(\alpha_1, \dots, \alpha_t)$ and $(\beta_1, \dots, \beta_t)$ be a sequence of elements from $\{1, \dots, r\}$ and $\{1, \dots, s\}$, respectively, such that (α_i, β_i) ’s are distinct. Then, for any $\theta = (q_1, \dots, q_t, q'_1, \dots, q'_r)$ (specifying the maximum number of queries for each permutation) we have*

$$\Delta_{\theta}^{\pm} \left((\bar{\Pi}_1, \dots, \bar{\Pi}_t, \hat{\Pi}); (\Pi_{\alpha_1}^{\oplus K_{\beta_1}}, \dots, \Pi_{\alpha_t}^{\oplus K_{\beta_t}}, \hat{\Pi}) \right) \leq \frac{\sigma}{2^n}$$

where $\sigma := 2 \sum_{\alpha=1}^r \left(\binom{\sigma_{\alpha}}{2} + \sigma_{\alpha} q'_{\alpha} \right)$ and $\sigma_{\alpha} = \sum_{\substack{1 \leq i \leq t \\ \alpha_i = \alpha}} q_i$ for every $1 \leq \alpha \leq r$.

Proof. The proof is similar to the proof of Lemma 2. Let $M^i, C^i \in (\{0, 1\}^n)^{q_i}$, $1 \leq i \leq t$, $X^{\alpha}, Y^{\alpha} \in (\{0, 1\}^n)^{q'_{\alpha}}$, $1 \leq \alpha \leq r$, be block-wise distinct tuples. From (6), we have that

$$\begin{aligned} I_{ideal} &= \Pr[\bar{\Pi}_i^{(q_i)}(M^i) = C^i, 1 \leq i \leq t, \Pi_{\alpha}^{(q'_{\alpha})}(X^{\alpha}) = Y^{\alpha}, 1 \leq \alpha \leq r] = \\ &= \prod_{i=1}^t \frac{1}{P(2^n, q_i)} \times \prod_{\alpha=1}^r \frac{1}{P(2^n, q'_{\alpha})}. \end{aligned}$$

We say that a tuple of keys (K_1, \dots, K_s) is “bad” if one of the following holds:

1. There are $1 \leq i, i' \leq t, 1 \leq j \leq q_i, 1 \leq j' \leq q_{i'}$ such that $(i, j) \neq (i', j')$, $\alpha_i = \alpha_{i'}$, and $K_{\beta_i} \oplus M_j^{\alpha_i} = K_{\beta_{i'}} \oplus M_{j'}^{\alpha_{i'}}$.
2. There are $1 \leq i \leq t, 1 \leq j \leq q_i, 1 \leq j' \leq q'_{\alpha_i}$ such that $K_{\beta_i} \oplus M_j^{\alpha_i} = X_{j'}^{\alpha_i}$.
3. There are $1 \leq i, i' \leq t, 1 \leq j \leq q_i, 1 \leq j' \leq q_{i'}$ such that $(i, j) \neq (i', j')$, $\alpha_i = \alpha_{i'}$, and $K_{\beta_i} \oplus C_j^{\alpha_i} = K_{\beta_{i'}} \oplus C_{j'}^{\alpha_{i'}}$.

4. There are $1 \leq i \leq t$, $1 \leq j \leq q_i$, $1 \leq j' \leq q'_{\alpha_i}$ such that $K_{\beta_i} \oplus C_j^{\alpha_i} = Y_{j'}^{\alpha_i}$.

Note that there are at most $\sum_{\alpha=1}^r \binom{\sigma_\alpha}{2}$ cases in the first and in the third items, and at most $\sum_{\alpha=1}^r \sigma_\alpha q'_\alpha$ cases in the second and fourth items.

If a key tuple is not bad, we say that it is a “good” key tuple. As in the proof of Lemma 2, for a good key tuple all the inputs (outputs) of each permutation are distinct. Thus, given a good tuple of keys (K_1, \dots, K_s) , it is easy to see that

$$\begin{aligned} \Pr[(\Pi_{\alpha_i}^{\oplus K_{\beta_i}})^{(q_i)}(M^i) = C^i, 1 \leq i \leq t, \Pi_{\alpha}^{(q'_\alpha)}(X^\alpha) = Y^\alpha, 1 \leq \alpha \leq r] = \\ = \prod_{\alpha=1}^r \frac{1}{P(2^n, \sigma_\alpha + q'_\alpha)} \geq I_{ideal}. \end{aligned}$$

It now remains to bound the probability that a random key tuple is bad. This can happen with one of the cases listed in items 1-4 where each case has probability 2^{-n} to occur. Hence, the probability that a random key tuple is bad, is at most $\frac{\sigma}{2^n}$, and the probability that a random key tuple is good is therefore at least $1 - \frac{\sigma}{2^n}$. The result follows by Theorem 1.

4.2 Main theorems

Theorem 4. Consider the BPEM cipher $BPEM[K_0, K_1, K_2; f_1, f_2]$ where the (secret) keys K_0, K_1, K_2 are selected uniformly and independently at random. Let q_* be the maximum number of queries to the encryption/decryption oracle, and let q_1, q_2 be the maximum numbers of queries to the public permutations f_1 and f_2 , respectively. Then,

$$\Delta_{BPEM}^{im-sprp}(q_*, q_1, q_2) \leq \frac{q_*(13q_* + 4q_1 + 4q_2)}{2^n}.$$

Proof. By Lemma 1, we know that our BPEM construction is same as

$$\text{LR}[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4}] \oplus (K'_6 * K'_5),$$

where K'_1, \dots, K'_6 are defined via (5) by K_1, K_2, K_3, K_4 . The matrix in (5) is lower triangular with non-zero diagonal, and hence non-singular. Thus, the “new” keys K'_1, \dots, K'_6 are also distributed uniformly and independently. As K'_5, K'_6 are independent of all the “ingredients” of $\text{LR}[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4}]$, it suffices to prove our result without the keys K'_5 and K'_6 .

Let Π_1, \dots, Π_4 be random permutations of $\{0, 1\}^n$ and let Π be a random permutation of $\{0, 1\}^{2n}$, all are independent of each other and independent of $\hat{f} = (f_1, f_2)$. Denote $\hat{F} = (f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_3}, f_2^{\oplus K'_4})$ and $\hat{\Pi} = (\Pi_1, \dots, \Pi_4)$. By Observation 2 and Lemma 3, we have⁷

$$\begin{aligned} \Delta_{q_*, q_1, q_2}^{\pm} \left((\text{LR}[\hat{F}], \hat{f}); (\text{LR}[\hat{\Pi}], \hat{f}) \right) &\leq \\ &\leq \Delta_{q_*, q_*, q_*, q_*, q_1, q_2}^{\pm} \left((\hat{F}, \hat{f}); (\hat{\Pi}, \hat{f}) \right) \leq \frac{4q_F(2q_F + q_1 + q_2)}{2^n}. \end{aligned}$$

⁷ Note that each query to the oracle construction $\text{LR}[g_1, g_2, g_3, g_4]$ translates to four queries - one to each permutation g_i , $i = 1, \dots, 4$

Finally, by applying the triangle inequality, Observation 1 and Theorem 2, the SPRP-advantage in the ideal model is

$$\begin{aligned}
\Delta_{q_*, q_1, q_2}^{\pm} \left((\text{LR}[\hat{F}], \hat{f}); (\Pi, \hat{f}) \right) &\leq \\
&\leq \Delta_{q_*, q_1, q_2}^{\pm} \left((\text{LR}[\hat{F}], \hat{f}); (\text{LR}[\hat{\Pi}], \hat{f}) \right) + \Delta_{q_*, q_1, q_2}^{\pm} \left((\text{LR}[\hat{\Pi}], \hat{f}); (\Pi, \hat{f}) \right) \leq \\
&\leq \frac{4q_F(2q_F + q_1 + q_2)}{2^n} + \Delta_{q_*}^{\pm} \left(\text{LR}[\hat{\Pi}]; \Pi \right) \leq \\
&\leq \frac{4q_*(2q_* + q_1 + q_2)}{2^n} + \frac{5q_*^2}{2^n} = \frac{q_*(13q_* + 4q_1 + 4q_2)}{2^n}.
\end{aligned}$$

The same argument can be used to show a similar bound for the single permutation 2-rounds BPPEM cipher.

Theorem 5. *Consider the single permutation BPPEM cipher $\text{BPPEM}[K_0, K_1, K_2; f, f]$ where the (secret) keys K_0, K_1, K_2 are selected uniformly and independently at random. Let q_* be the maximum number of queries to the encryption/decryption oracle, and let q be the maximum number of queries to the public permutation f . Then,*

$$\Delta_{\text{BPPEM}[K_0, K_1, K_2; f, f]}^{\text{im-sprp}}(q_*, q) \leq \frac{q_*(21q_* + 8q)}{2^n}$$

Remark 4. The difference in the bounds we received in Theorems 4 and 5 are due only to the difference in the value of σ in the application of Lemma 3.

We also comment that the same bounds hold in the single key variants. By (5) we have

$$\begin{aligned}
\text{BPPEM}[K, K, K; f_1, f_2] &= \text{LR}[f_1^{\oplus K'_1}, f_1^{\oplus K'_2}, f_2^{\oplus K'_2}, f_2^{\oplus K'_3}], \\
\text{BPPEM}[K, K, K; f, f] &= \text{LR}[f^{\oplus K'_1}, f^{\oplus K'_2}, f^{\oplus K'_2}, f^{\oplus K'_3}]
\end{aligned}$$

where

$$\begin{pmatrix} K'_1 \\ K'_2 \\ K'_3 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} K_R \\ K_L \end{pmatrix}.$$

For both constructions, the “new” keys K'_1, K'_2, K'_3 are no longer independent, so we need to generalize lemma 3 as stated below.

Lemma 4. *Let $\Pi_1, \dots, \Pi_r, \bar{\Pi}_1, \dots, \bar{\Pi}_t$ be independent random permutations of $\{0, 1\}^n$ and K_1, \dots, K_s be chosen uniformly and independently from $\{0, 1\}^n$. We write $\hat{\Pi}$ to denote (Π_1, \dots, Π_r) . Let $(\alpha_1, \dots, \alpha_t)$ be a sequence of elements from $\{1, \dots, r\}$. Let M be a binary matrix of size $t \times s$, with no zero rows, satisfying the following condition: for every $1 \leq i_1 < i_2 \leq t$ such that $\alpha_{i_1} = \alpha_{i_2}$, the i_1^{th} and i_2^{th} rows of M are distinct. Let $K'_i := \sum_{j=1}^s M_{ij} K_j$, for every $1 \leq i \leq t$.*

Then, for any $\theta = (q_1, \dots, q_t, q'_1, \dots, q'_r)$ (specifying the maximum number of queries) we have

$$\Delta_{\theta}^{\pm} \left((\bar{\Pi}_1, \dots, \bar{\Pi}_t, \hat{\Pi}); (\Pi_{\alpha_1}^{\oplus K'_{\beta_1}}, \dots, \Pi_{\alpha_t}^{\oplus K'_{\beta_t}}, \hat{\Pi}) \right) \leq \frac{\sigma}{2^n}$$

where $\sigma := 2 \sum_{\alpha=1}^r \left(\binom{\sigma_{\alpha}}{2} + \sigma_{\alpha} q'_{\alpha} \right)$ and σ is as defined in Lemma 3.

We skip the proof of this lemma as it is similar to that of Lemma 3. Similarly to the proof of Theorem 4 (while using Lemma 4 instead of Lemma 3), we can obtain the following bound.

Theorem 6. Consider the single key BPEM cipher $\text{BPEM}[K, K, K; f_1, f_2]$ where the (secret) key K is selected uniformly at random. Let q_* be the maximum number of queries to the encryption/decryption oracle, and let q_1, q_2 be the maximum numbers of queries to the public permutations f_1 and f_2 , respectively. Then,

$$\Delta_{\text{BPEM}[K, K, K; f_1, f_2]}^{\text{im-sprp}}(q_*, q_1, q_2) \leq \frac{q_*(13q_* + 4q_1 + 4q_2)}{2^n}.$$

Finally, similarly to the proof of Theorem 5 (while using Lemma 4 instead of Lemma 3, and using Theorem 3 instead of Theorem 2), we obtain the following bound.

Theorem 7. Consider the single key single permutation BPEM cipher $\text{BPEM}[K, K, K; f, f]$ where the (secret) key K is selected uniformly at random. Let q_* be the maximum number of queries to the encryption/decryption oracle, and let q be the maximum number of queries to the public permutation f . Then,

$$\Delta_{\text{BPEM}[K, K, K; f, f]}^{\text{im-sprp}}(q_*, q) \leq \frac{q_*(16q_* + 8q)}{2^n} + \frac{17q_*^2}{2^n - 1} + \frac{q_*^2}{2^{2n}}.$$

5 A distinguishing attack on BPEM

In this section we describe a distinguishing attack on BPEM that uses $O(2^{n/2})$ queries. This is the same attack as the one described in [20, Section 3.2] for the 4-rounds Luby-Rackoff with internal permutations, not at all surprising, since we showed (in Section 2.3) that BPEM can be viewed as a 4-rounds Luby-Rackoff with internal (keyed) permutations. Nevertheless, for the sake of completeness, we describe and analyze the attack in this BPEM terminology. We will use the following technical lemma.

Lemma 5. If $x, y, \rho \in \{0, 1\}^n$ such that

$$\begin{aligned} x \oplus (\text{BPEM}[K_0, K_1, K_2; f_1, f_2](x * \rho))_L &= \\ = y \oplus (\text{BPEM}[K_0, K_1, K_2; f_1, f_2](y * \rho))_L \end{aligned} \quad (7)$$

then $x = y$.

Proof. Denote

$$\begin{aligned}\check{x} &:= \text{LR}^2[f_1]((x * \rho) \oplus K_0) \oplus K_1, \\ \check{y} &:= \text{LR}^2[f_1]((y * \rho) \oplus K_0) \oplus K_1.\end{aligned}$$

By (4) and (1) we have that

$$\begin{aligned}\text{BPEM}[K_0, K_1, K_2; f_1, f_2](x * \rho) &= \text{EM}_{K_0, K_1, K_2}^{\text{LR}^2[f_1], \text{LR}^2[f_2]}(x * \rho) = \\ &= \text{LR}^2[f_2](\text{LR}^2[f_1]((x * \rho) \oplus K_0) \oplus K_1) \oplus K_2 = \text{LR}^2[f_2](\check{x}) \oplus K_2,\end{aligned}$$

hence, by (3),

$$\begin{aligned}(\text{BPEM}[K_0, K_1, K_2; f_1, f_2](x * \rho))_L &= (\text{LR}^2[f_2](\check{x}) \oplus K_2)_L = \\ &= \check{x}_L \oplus f_2(\check{x}_R) \oplus (K_2)_L = x \oplus (K_0)_L \oplus (K_1)_L \oplus f_1(\rho \oplus (K_0)_R) \oplus f_2(\check{x}_R) \oplus (K_2)_L.\end{aligned}$$

Similarly

$$\begin{aligned}(\text{BPEM}[K_0, K_1, K_2; f_1, f_2](y * \rho))_L &= \\ &= y \oplus (K_0)_L \oplus (K_1)_L \oplus f_1(\rho \oplus (K_0)_R) \oplus f_2(\check{y}_R) \oplus (K_2)_L.\end{aligned}$$

Therefore we get from (5) that $f_2(\check{x}_R) = f_2(\check{y}_R)$, hence, since f_2 is injective, $\check{x}_R = \check{y}_R$. Therefore, using (3) again,

$$\begin{aligned}\rho \oplus (K_0)_R \oplus f_1(x \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R)) \oplus (K_1)_R &= \\ = \rho \oplus (K_0)_R \oplus f_1(y \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R)) \oplus (K_1)_R,\end{aligned}$$

hence

$$f_1(x \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R)) = f_1(y \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R)).$$

Since f_1 is injective we get that

$$x \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R) = y \oplus (K_0)_L \oplus f_1(\rho \oplus (K_0)_R),$$

hence $x = y$.

Proposition 2. *Consider the BPEM cipher $\text{BPEM}[K_0, K_1, K_2; f_1, f_2]$ with arbitrary (secret) keys K_0, K_1, K_2 . Let q_* be the maximum number of queries to the encryption oracle. Then,*

$$\Delta_{\text{BPEM}}^{\text{prp}}(q_*) \geq 1 - e^{-\frac{q_*(q_*-1)}{2(2^n+1)}}.$$

Remark 5. Note that Proposition 2 implies that the adversary advantage becomes non-negligible for $q_* = \Omega(2^{n/2})$.

Proof. Fix an n -bit string ρ and q_* distinct n -bit strings $\omega_1, \omega_2, \dots, \omega_{q_*}$. We query the encryption oracle for the plaintexts $\omega_1 * \rho, \omega_2 * \rho, \dots, \omega_{q_*} * \rho$, and let $\sigma_1, \sigma_2, \dots, \sigma_{q_*}$ be the corresponding ciphertexts. We now search for collisions between the q_* n -bit strings $\omega_1 \oplus (\sigma_1)_L, \omega_2 \oplus (\sigma_2)_L, \dots, \omega_{q_*} \oplus (\sigma_{q_*})_L$. By Lemma 5 there will be no collision if the oracle encrypts using $\text{BPEM}[K_0, K_1, K_2; f_1, f_2]$. By contrast, if the oracle encrypts by applying a randomly chosen permutation of $\{0, 1\}^{2n}$ then the probability there is no collision is at most

$$\prod_{k=1}^{q_*-1} \left(1 - \frac{k(2^n - 1)}{2^{2n} - k}\right) \leq \prod_{k=1}^{q_*-1} \left(1 - \frac{k}{2^n + 1}\right) \leq \prod_{k=1}^{q_*-1} e^{-\frac{k}{2^n + 1}} = e^{-\frac{q_*(q_*-1)}{2(2^n + 1)}}.$$

6 A practical constructions of a 256-bit cipher

In this section, we demonstrate a practical construction of a 256-bit block cipher based on the 2-rounds BPEM cipher, where the underlying permutation is AES.

Definition 4 (EM256AES: a 256-bit block cipher). *Let ℓ_1 and ℓ_2 be two 128-bit keys and let K_0, K_1, K_2 be three 256-bit secret keys (assume $\ell_1, \ell_2, K_0, K_1, K_2$ are selected uniformly and independently at random). Let the permutations f_1 and f_2 be the AES encryption using ℓ_1 and ℓ_2 as the AES key, respectively. The 256-bit block cipher EM256AES is defined as the associated instantiation of the 2-rounds BPEM cipher $\text{BPEM}[K_0, K_1, K_2; f_1, f_2]$. Usage of EM256AES:*

- ℓ_1 and ℓ_2 are determined during the setup phase, and can be made public (e.g., sent from the sender to the receiver as an IV).
- K_0, K_1, K_2 are selected per encryption session.

The single key EM256AES is the special case where a single value $K \in \{0, 1\}^{256}$ and a single value $\ell \in \{0, 1\}^{128}$ are selected uniformly and independently at random, and the EM256AES cipher uses $K_0 = K_1 = K_2 = K$ and $\ell_1 = \ell_2 = \ell$.

Hereafter, we use the single key EM256AES. To establish security properties for EM256AES, we make the standard assumption about AES with a secret key that is selected (uniformly at random): an adversary has negligible advantage in distinguishing AES from a random permutation of $\{0, 1\}^{128}$ even after seeing a (very) large number of plaintext-ciphertext pairs (i.e., the assumption is that AES satisfies its design goals ([1], Section 4). This assumption is certainly reasonable if the number of blocks that are encrypted with the same keys is limited to be much smaller than 2^{64} . Therefore, in our context, we can consider assigning the randomly selected key ℓ at setup time, as an approximation for a random selection of the permutation $f_1 = f_2$. Under this assumption, we can rely on the result of Theorem 7 for the security of EM256AES.

***EM256AES* efficiency:** An encryption session between two parties requires exchanging a 256-bit secret key and transmitting a 128-bit IV ($= \ell$). One key (and IV) can be used for N blocks as long as we keep $N \ll 2^{64}$.

Computing one (256-bit) ciphertext involves 4 AES computations (with the IV as the AES key) plus a few much cheaper XOR operations. Let us assume that the encryption is executed on a platform that has the capability of computing AES at some level of performance. If the *EM256AES* encryption (decryption) is done in a serial mode, we can estimate the encryption rate to be roughly half the rate of AES (serial) computation on that platform (4 AES operations per one 256-bit block). Similarly, if the *EM256AES* encryption is done in a parallelized mode, we can estimate the throughput to be roughly half the throughput of AES.

***EM256AES* performance:** To test the actual performance of *EM256AES*, and validate our predictions, we coded an optimized implementation of *EM256AES*. Its performance is reported here.

The performance was measured on an Intel Core i7-4700MQ (microarchitecture Codename Haswell) where the enhancements (Intel Turbo Boost Technology, Intel Hyper-Threading Technology, and Enhanced Intel Speedstep Technology) were disabled. The code used the AES instructions (AES-NI) that are available on such modern processors.

On this platform, we point out the following baseline: the performance of AES (128-bit key) in a parallelized mode (CTR) is 0.63 C/B, and in a serial mode (CBC) it is 4.44 cycles per byte (C/B hereafter).

The measured performance of our *EM256AES* implementation was 1.44 C/B for the parallel mode, and 8.92 C/B for the serial mode. The measured performance clearly matches the predictions.

It is also interesting to compare the performance of *EM256AES* to another 256-bit cipher. To this end, we prepared an implementation of Rijndael256 cipher [6]². For details on how to code Rijndael256 with AES-NI, see [10]). Rijndael256 (in ECB mode) turned out to be much slower than *EM256AES*, performing at 3.85 C/B.

7 Discussion

In this work, we showed how to construct a large family of balanced permutations, and analyzed the resulting new variation, BPEM, of the EM cipher.

The resulting $2n$ -bit block cipher is obtained by using a permutation of $\{0, 1\}^n$ as a primitive. The computational cost of encrypting (decrypting) one $2n$ -bit block is 4 evaluations of a permutation of $\{0, 1\}^n$ (plus a relatively small overhead). Note that this makes BPEM readily useful in practice, for example to define a 256-bit cipher, because “good” permutations of $\{0, 1\}^{128}$ are available.

² AES is based on the Rijndael block cipher. While AES standardizes only a 128 block size, the Rijndael definitions support both 128-bit and 256-bit blocks

We demonstrated the specific cipher *EM256AES*, which is based on AES, and showed that its throughput is (only) half the throughput of AES (and 2.5 times faster than Rijndael256).

A variation on the way by which BPEM can be used, would make it a tweakable $2n$ -bit block cipher. Here, the public IV ($=\ell$) can be associated with each encrypted block as an identifier, to be viewed as the tweak. The implementation would switch this tweak for each block. To randomize the keys for the (public) permutations, an additional encryption (using some secret key) is necessary.

The expression of the advantage in Theorem 4 behaves linearly with the number of queries to the public permutations, and quadratically with the number of queries to the encryption/decryption oracle. This reflects the intuition that the essential limitations on the number of adversary queries should be on the encryption/decryption invocations, while weaker (or perhaps no) limitations should be imposed on the number of queries to the public permutations. It also suggests the following protocol, where the secret keys are changed more frequently than the random permutations. *Choose the public permutations for a period of, say, $\frac{1}{1000}2^{2n/3}$ blocks, divided into $2^{n/3}$ sessions of $\frac{1}{1000}2^{n/3}$ blocks. Change the secret keys every session.* This way, the relevant information on the block cipher, from a specific choice of keys, is limited to a session, while the adversary can accumulate relevant information from replies to the public permutations across sessions. Therefore, q_* is limited to $\frac{1}{1000}2^{n/3}$, while $q_* + q_1 + q_2$ is limited to $\frac{1}{1000}2^{2n/3}$. Theorem 4 guarantees that this usage is secure.

References

1. –, Announcing request for candidate algorithm nominations for the Advanced Encryption Standard (AES), <http://csrc.nist.gov/CryptoToolkit/aes/pre-round1/aes.9709.htm> (1997).
2. A. Bogdanov, L. R. Knudsen, G. Leander, F-X Standaert, J. P. Steinberger, and E. Tischhauser, Key-alternating ciphers in a provable setting: encryption using a small number of public permutations (extended abstract), in *Advances in cryptology—EUROCRYPT 2012*, Lecture Notes in Comput. Sci., 7237, Springer, Heidelberg, 45–62 (2012).
3. S. Chen and J. Steinberger, Tight security bounds for key-alternating ciphers, in *Advances in cryptology—EUROCRYPT 2014*, 327–350, Lecture Notes in Comput. Sci., 8441, Springer, Heidelberg, 2014.
4. J.-S. Coron, J. Patarin, Y. Seurin, The random oracle model and the ideal cipher model are equivalent, in *Advances in cryptology—CRYPTO 2008*, 1–20, Lecture Notes in Comput. Sci., 5157, Springer, Berlin.
5. J. Daemen, Limitations of the Even-Mansour construction, in *Advances in cryptology—ASIACRYPT ’91*, Lecture Notes in Computer Science, 739, Springer, Berlin, (H. Imai, R. L. Rivest, T. Matsumoto, editors) 495–498 (1993).
6. J. Daemen, V. Rijmen, AES Proposal: Rijndael (National Institute of Standards and Technology), <http://csrc.nist.gov/archive/aes/rijndael/Rijndael-ammended.pdf> (1999).
7. I. Dinur, O. Dunkelman, N. Keller, A. Shamir, Key Recovery Attacks on 3-round Even-Mansour, 8-step LED-128, and full AES², <http://eprint.iacr.org/2013/391> (2013).

8. S. Even, Y. Mansour, A construction of a cipher from a single pseudorandom permutation, *J. Cryptology* **10** (1997), no. 3, 151–161.
9. C. Gentry, Z. Ramzan, Eliminating random permutation oracles in the Even-Mansour cipher, in *Advances in cryptology—ASIACRYPT 2004*, 32–47, Lecture Notes in Comput. Sci., 3329, Springer, Berlin.
10. S. Gueron, Intel Advanced Encryption Standard (AES) Instructions Set (Rev 3.01), <http://software.intel.com/sites/default/files/article/165683/aes-wp-2012-09-22-v01.pdf> (2014).
11. R. Lampe, Y. Seurin, Security Analysis of Key-Alternating Feistel Ciphers, in *Fast Software Encryption — FSE 2014*, 243–264, Lecture Notes in Comput. Sci., 8540, (2014).
12. R. Lampe, J. Patarin, Y. Seurin, An Asymptotically Tight Security Analysis of the Iterated Even-Mansour Cipher, in *Advances in cryptology—ASIACRYPT 2012*, 278–295, Lecture Notes in Comput. Sci., 7658, (2012).
13. U. Maurer, R. Renner, C. Holenstein, Indifferentiability, impossibility results on reductions, and applications to the random oracle methodology, in *Theory of cryptography*, 21–39, Lecture Notes in Comput. Sci., 2951, Springer, Berlin.
14. M. Nandi, The characterization of Luby-Rackoff and its optimum single-key variants, in *Progress in cryptology—INDOCRYPT 2010*, 82–97, Lecture Notes in Comput. Sci., 6498, Springer, Berlin.
15. I. Nikolić, L. Wang, S. Wu, Cryptanalysis of Round-Reduced LED. In *FSE*, 2013. To appear in Lecture Notes in Computer Science.
16. J. Patarin, *Étude des générateurs de permutations pseudo-aléatoires basés sur le schéma du D.E.S.*, INRIA, Rocquencourt, 1991.
17. J. Patarin, Luby-Rackoff: 7 rounds are enough for $2^{n(1-\varepsilon)}$ security, in *Advances in cryptology—CRYPTO 2003*, 513–529, Lecture Notes in Comput. Sci., 2729, Springer, Berlin.
18. G. Piret, Luby-Rackoff revisited: on the use of permutations as inner functions of a Feistel scheme, *Des. Codes Cryptogr.* **39** (2006), no. 2, 233–245.
19. J. P. Steinberger, Improved Security Bounds for Key-Alternating Ciphers via Hellinger Distance, IACR Cryptology ePrint Archive 2012: 481 (2012).
20. J. Treger, J. Patarin, Generic attacks on Feistel networks with internal permutations, in *Progress in cryptology—AFRICACRYPT 2009*, 41–59, Lecture Notes in Comput. Sci., 5580, Springer, Berlin.